

WEST Search History

[Hide Items](#) [Restore](#) [Clear](#) [Cancel](#)

DATE: Wednesday, March 23, 2005

<u>Hide?</u>	<u>Set Name</u>	<u>Query</u>	<u>Hit Count</u>
<i>DB=USPT,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR</i>			
<input type="checkbox"/>	L12	l2 and usb	5
<input type="checkbox"/>	L11	L8 same tree	3
<input type="checkbox"/>	L10	L8 same tree same device	0
<input type="checkbox"/>	L9	L8 same (root or hub)	11
<input type="checkbox"/>	L8	(power\$4 near2 (down or off)) near5 (separat\$4 or independent\$2 or individual\$2)	2064
<input type="checkbox"/>	L7	l3 not vehicle	12
<input type="checkbox"/>	L6	l3 same maintain\$4	1
<input type="checkbox"/>	L5	l3 same request\$4	0
<input type="checkbox"/>	L4	L2 with (root or hub)	23
<input type="checkbox"/>	L3	L2 same (root or hub)	34
<input type="checkbox"/>	L2	(select\$4 or independent\$2 or individual\$4) near3 (suspen\$5 or idl\$4 or inactiv\$4) near3 (device or terminal or unit or equipment)	1585
<input type="checkbox"/>	L1	(select\$4 or independent or individual\$4) near3 (suspen\$5 or idl\$4 or inactiv\$4) near3 (device or terminal or unit or equipment)	1430

END OF SEARCH HISTORY

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#) [Generate Collection](#) [Print](#)

L12: Entry 3 of 5

File: USPT

Dec 11, 2001

DOCUMENT-IDENTIFIER: US 6330639 B1

** See image for Certificate of Correction **

TITLE: Method and apparatus for dynamically changing the sizes of pools that control the power consumption levels of memory devices

Brief Summary Text (5):

The system constraints and tradeoffs described above with respect to computer devices in general apply equally to memory devices in memory system. In their active or most power-hungry mode, memory devices such as dynamic random access memory (DRAM) devices operate faster than they do when they are in inactive mode (e.g., standby or nap mode). However, DRAM devices in their active mode also consume far more power than they do when they are in inactive mode. As a result, to maintain a balance between performance and power consumption (and heat dissipation), some fixed number of DRAM devices may need to be kept in an inactive mode to conserve power and reduce heat dissipation. The number of devices in active mode and the number of devices in inactive mode can be specified by the Basic Input/Output Program (BIOS) at system start up (boot) or system reset. Management of which devices are in active mode and which devices are in inactive mode can be accomplished through a definition of pools of devices that are used to keep track of the operational mode or power state (e.g., active or inactive) of the individual devices. A pool of devices in this context refers to a mapping or list of devices that are in a specific operational mode or power state. For example, one pool may be maintained to keep track of the devices that are in active mode and another pool may be maintained to keep track of the devices that are in inactive mode. Under such a power management scheme, the devices represented in one of the pools are assumed to be operating in a certain operational mode or power state and therefore consuming a certain amount of power. For example, devices that are represented in the active pool are assumed to be operating in active mode. The number of devices in each pool may be examined to determine the amount of power being used by the entire memory system. The different pools utilized to keep track of the operational mode or power state of the various memory devices are also referred to as the power-control or power-saving pools hereinafter.

Detailed Description Text (4):

FIG. 1 shows a block diagram of one embodiment of a system configuration in which the teachings of the present invention are implemented. The system configuration 100 includes a plurality of central processing units (CPUs) 101a-d, a memory control hub (also referred to as memory control unit) 111, a P64 control unit 121, an Input/Output (IO) control unit 131, a graphics controller 141 coupled to a graphics subsystem 151, and a plurality of memory devices 161. For the purposes of the present specification, the term "processor" or "CPU" refers to any machine that is capable of executing a sequence of instructions and shall be taken to include, but not be limited to, general-purpose microprocessors, special purpose microprocessors, multi-media controllers and microcontrollers, etc. In one embodiment, the CPUs 101a-101d are general-purpose microprocessors that are capable of executing an Intel Architecture instruction set. The CPUs 101a-101d, the P64 control unit 121, the IO control unit 131, and the AGP graphics control unit 141 access the system memory devices 161 via the memory control unit 111. The memory control unit 111, in one embodiment, is responsible for servicing all memory transactions that target the system memory devices 161. The memory control unit 111

can be a stand-alone unit, an integrated part of a chipset, or a part of some larger unit that control the interfaces between various system components and the system memory devices 161. The P64 control unit 121 provides the interface control between a plurality of PCI-64 slots 125 and the memory control unit 111. The IO control unit 131 provides the interface control between the memory unit 111 and various IO devices and ports including the PCI slots and PCI agents 133, a plurality of USB ports 135, a plurality of IDE ports 137, and other IO devices 139. The AGP graphics control unit 141 provides the interface control between the graphics subsystem 151 and the memory control unit 111. The structure and functions of the memory control unit 111 are described in greater details below.

[Previous Doc](#)[Next Doc](#)[Go to Doc#](#)

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#) [Generate Collection](#) [Print](#)

L12: Entry 4 of 5

File: USPT

Mar 6, 2001

DOCUMENT-IDENTIFIER: US 6199134 B1

TITLE: Computer system with bridge logic that asserts a system management interrupt signal when an address is made to a trapped address and which also completes the cycle to the target address

Abstract Text (1):

A computer system includes a South bridge logic that connects an expansion bus to one or more secondary expansion busses and peripheral devices. The South bridge logic includes internal control devices that are targets for masters on the expansion bus. The target devices couple to the expansion bus through a common expansion target interface, which monitors and translates master cycles on the expansion bus on behalf of the target devices. The South bridge includes an ACPI/power management logic capable of supporting a Device Idle mode in which selected I/O device may be placed in a low power state. To prevent cycles from being run to a device in a low power state, the ACPI/power management includes status registers that are used to determine when a device in low power mode is the target of an expansion bus cycle. If such a cycle occurs, the cycle is intercepted and an SMI signal is transmitted to the CPU. In addition, the target interface responds to the master by asserting a retry signal. When the transaction is retried, the cycle is passed to the target, which responds with an invalid data signal. The CPU by this time, or at some subsequent time realizes that the target was asleep based upon processing of the SMI signal. The CPU then either re-executes the cycle when the device is removed from the low power state, or else simply rejects the invalid data.

Detailed Description Text (8):

Referring still to FIG. 2, another bridge logic device 100 also preferably connects to expansion bus 55. This bridge logic device 100 (which sometimes is referred to a "South bridge") couples or "bridges" the primary expansion bus 55 to other secondary expansion busses. These other secondary expansion busses may include an ISA (Industry Standard Architecture) bus, a sub-ISA bus, a USB (Universal Serial Bus), an IDE bus, an IEEE 1394 standard (or "fire wire") bus, or any of a variety of other busses that are available or may become available in the future. In the preferred embodiment of FIG. 2, the South bridge logic 100 couples to a sub-ISA bus 87, a USB bus 97 and an IDE bus 98. As will be apparent to one skilled in this art, various peripheral devices may be coupled to each of these busses. Thus, as shown in the preferred embodiment of FIG. 2, an MSIO (Mobile Super I/O) chip 92 connects to the sub-ISA bus, as does an audio card 85, modem 90 and BIOS ROM 91. Similarly, according to the preferred embodiment, a multi-bay configuration couples to the IDE bus 98 and to the MSIO 92. The multi-bay configuration preferably includes three independent bays connected in any master/slave, primary/secondary configuration supporting hot-plugging of IDE devices in the bays. As one skilled in the art will understand various IDE devices are available including hard drives, CD ROM drives, floppy drives, tape drives, and the like. The USB 97 supports various peripherals, especially video peripherals such as video cameras for teleconferencing purposes. In addition to the busses discussed above, the South bridge also preferably connects to interrupt signal lines, power management signal lines, and a MSB (MSIO Serial Bus).

Detailed Description Text (9):

As noted above, the computer system 10 preferably comprises a laptop computer. To facilitate use of the computer system 10 within a traditional desktop environment, an expansion base preferably is available for connecting the laptop computer system to additional peripheral devices and to a computer network via an Ethernet bus. The expansion base may include any type of configuration, including a port replicator or docking station. The present invention envisions, however, that the expansion base also includes a South bridge logic (not shown specifically) for coupling components on the PCI bus to other components coupled to the expansion base. Thus, as shown in FIG. 2, a South bridge logic device within the expansion base 70 couples the PCI bus to various expansion busses including an ISA bus 81, USB 77 and IDE bus 79. The following discussion focuses on the architecture of the South bridge logic 100. This same architecture may be used (with some minor modifications), if desired, for the South bridge logic within the expansion base 70. Alternatively, the present invention also may be implemented by using a prior art South bridge device in the expansion base 70.

Detailed Description Text (20):

Referring still to FIG. 3, the address decoder 190 couples to the PCI target interface 185 to indicate if the PCI signals on the PCI bus comprise valid PCI signals for decoding by address decoder 190. The address decoder 190 indicates to the target interface 185 if the PCI address value matches an address located in the South bridge or on one of the secondary expansion busses coupled to the South bridge, such as the IDE bus of some other bus such as a USB bus. The address decoder 190 also produces a number of chip select output signals (not specifically shown) that are provided to each of the internal IMAX targets and which are used to select one of the internal targets as the target for an IMAX cycle. Stated differently, the chip select signals identify which of the IMAX internal targets are being addressed by the external PCI master. The address decoder 190 is closely linked to the configuration registers. The configuration registers 195 preferably include an Internal I/O Positive Decode Enable Configuration Register that, among other things, indicates if the South bridge 100 is responsible for decoding I/O cycles to an internal South bridge register. In the preferred embodiment, the appropriate bit is set in this register causing the address decoder 190 to decode I/O cycles to internal South bridge devices, while the same register bit is reset in the South bridge in the expansion base 70 (FIG. 2). Thus, in the preferred embodiment, the South bridge 100 exclusively decodes cycles to unique internal I/O register addresses.

CLAIMS:

22. A computer system as in claim 20, wherein said at least one secondary expansion bus comprises one or more of the following busses: an ISA bus; a USB bus; or an IDE bus.

[Previous Doc](#)

[Next Doc](#)

[Go to Doc#](#)

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#) [Generate Collection](#) [Print](#)

L12: Entry 2 of 5

File: USPT

Dec 17, 2002

DOCUMENT-IDENTIFIER: US 6496938 B1

TITLE: Enhanced PCI clock control architecture

Abstract Text (1):

A clock control technique allows reducing the power consumption of devices connected to a computer bus. Individual idle devices can be disconnected from the bus clock by a device clock controller and placed in a low-power state without waiting for all devices on the bus to go idle. When individual devices are idle, transactions on the bus are monitored and unclaimed transactions are claimed by the device clock controller, which then forces a retry of the transaction and reconnects the clock to the idle devices. This brings these devices from the low-power state to a full power state, where they are capable of claiming the transaction when it is retried.

Detailed Description Text (8):

In a disclosed embodiment, the Host Bridge 114 can be a 440LX North Bridge Integrated Circuit by Intel Corporation, also known as the PCI AGP Controller (PAC). The South Bridge 124 can be a PIIIX4, also by Intel Corporation. The Host Bridge 114 and the South Bridge 124 could be replaced with other bridge chip sets without departing from the spirit and scope of the invention. The Host Bridge 114 and the South Bridge 124 provide capabilities other than bridging between the processor socket 116 and the PCI bus 110, and between the PCI bus 110 and the ISA bus 112. Specifically, the disclosed Host Bridge 114 includes interface circuitry for the AGP connector 118, the memory subsystem 120, and the AGP 122. The disclosed ISA bridge 124 further includes an internal enhanced IDE controller for controlling up to four enhanced IDE drives 126, and a universal serial bus (USB) controller for controlling USB ports 128.

CLAIMS:

2. A computer bus clock control mechanism as in claim 1, wherein the clock controller deasserts a clock running signal to the device clock controller to indicate a request to stop the clock signal and asserts the clock running signal to the device clock controller to indicate the clock is running, wherein the device clock controller deasserts an individual copy of the clock running signal to each of the at least one other bus-connected devices responsive to the deassertion of the clock running signal by the clock controller, wherein each of the at least one other device that is not idle asserts the individual copy of the clock running signal responsive to the deassertion of the individual copy of the clock running signal by the device clock controller, wherein the device clock controller asserts the clock running signal responsive to the assertion of any of the individual copies of the clock running signal, and wherein the clock controller signals the clock generator to stop and start the clock signal responsive to the assertion and deassertion of the clock running signal from the device clock controller.

18. A computer system as in claim 16, wherein the clock controller deasserts a clock running signal to the device clock controller to indicate a request to stop the clock and asserts the clock running signal to indicate the clock is running, wherein the device clock controller deasserts an individual copy of the clock running signal to each of the at least one peripheral device responsive to the

deassertion of the clock running signal by the clock controller, wherein each of the at least one peripheral device that is not idle asserts the individual copy of the clock running signal responsive to the deassertion of the individual copy of the clock running signal by the device clock controller, wherein the device clock controller asserts the clock running signal responsive to the assertion of any of the individual copies of the clock running signal, and wherein the clock controller signals the clock generator to stop and start the clock signal responsive to the assertion and deassertion of the clock running signal from the device clock controller.

[Previous Doc](#)[Next Doc](#)[Go to Doc#](#)

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#) [Generate Collection](#) [Print](#)

L9: Entry 1 of 11

File: USPT

Aug 24, 2004

DOCUMENT-IDENTIFIER: US 6782472 B2

TITLE: Method of initializing a memory controller by executing software in a second memory to wake up a system

Detailed Description Text (10):

In step 210, the processor transitions to the sleep state. One way to accomplish this transition is to set the appropriate bits in a sleep enable register. Either a software or hardware process then detects that this bit is set and asserts a sleep signal to the appropriate components. Processor clock 50 is powered down. Powering down may be accomplished by disconnecting power from the device itself, or may be accomplished by electrically disconnecting the incoming signal from the internal distribution lines internal to each chip. For example, processor clock 50 may be left running but the processor may electrically disconnect the incoming clock signal so the processor's internal components are not being clocked. Likewise, individual devices may be powered down with circuitry internal to the devices that prevent the flow of power to some or all of the components inside the device. In an RDRAM system, memory controller hub 20, main memory 30, and memory clock 40 are powered down. When the main memory is power down, its contents are not lost, but the main memory devices transition to a power down state that consumes very little power. An internal self refresh mechanism within main memory 20 keeps the memory contents when main memory is powered down. Also, memory clock 40 transitions to a low power state. In the low power state, physical power may or may not be removed.

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#) [Generate Collection](#) [Print](#)

L9: Entry 3 of 11

File: USPT

Sep 14, 1999

DOCUMENT-IDENTIFIER: US 5953174 A

TITLE: Magnetic tape drive unit having reduced power consumption

Brief Summary Text (9):

In accordance with an aspect of the present invention, there is provided a compact magnetic tape drive unit comprising a write circuit for modulating given data and then writing the data modulated through a magnetic head onto a magnetic tape; a read circuit for reading contents written on the magnetic tape by means of the magnetic head and then demodulating the contents read; a cartridge load circuit for controlling a load operation that a magnetic tape cartridge put into the drive unit is moved to a position where a hub of the magnetic tape cartridge is rotatable, and also controlling an unload operation reverse to the load operation; a tape thread circuit for controlling a thread operation that the magnetic tape is drawn from the magnetic tape cartridge loaded and is then wound around a machine reel, and also controlling an unthread operation reverse to the thread operation; a power supply for supplying power to each of the circuits; power on/off means for switching on or off supply of power from the power supply to each of the circuits individually; and power supply control means for monitoring an operating condition of each of the circuits and controlling the power on/off means according to whether the supply of power to each of the circuits is required or not.

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#) [Generate Collection](#) [Print](#)

L4: Entry 9 of 23

File: USPT

Dec 1, 1987

DOCUMENT-IDENTIFIER: US 4710916 A

TITLE: Switching apparatus for burst-switching communications system

Detailed Description Text (192):

In summary, a message burst passes through the hub switch from an incoming channel of one link group to an outgoing channel of another link group in the following manner. Bytes of the burst arriving at the origin hub switching unit 501 on an incoming link channel are buffered in the switching unit memory 516. The first bytes, or header, of the burst contain address information; one byte, specifically the second byte, designating the destination link group and therefore the destination switching unit. The received bytes are queued for transmission on the hub bus. A hub channel in which the origin switching unit is transmit idle and the destination switching unit is receive idle is selected. The bytes of the burst are loaded onto the selected hub channel, one byte during each hub channel frame. A byte is transferred directly between the hub switch elements of adjacent intervening switching units on each clock tick without passing through the memories 516. Upon arriving at the destination switching unit, each byte is stored in the memory. The header bytes are interpreted to determine the appropriate output link group, if more than one link group is associated with the destination switching unit. The bytes are queued on the appropriate outbound link, and output begins on the first idle outbound link channel.

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#) [Generate Collection](#) [Print](#)

L7: Entry 1 of 12

File: USPT

Jan 3, 1989

DOCUMENT-IDENTIFIER: US 4796227 A
** See image for Certificate of Correction **
TITLE: Computer memory system

Detailed Description Text (58):

A decoding circuit according to the present invention avoids the long delays associated with prior art static decoding circuitry while retaining the desirable features of a static decoding system. Basically, the decoding circuit of the present invention replaces the series connected pullup transistors of the NOR gates used in prior art decoding circuits with a tree-like structure which provides one connection path to the active power terminal. The active power terminal is at the root node of said tree. This root node connects this power rail to two or more branches, each of which may in turn be connected to a root node of a "sub-tree". In a decoding circuit for decoding an N-bit binary address, the branching ends at $2^{\sup{N}}$ "leaf" nodes. The word lines are connected to these leaf nodes. With the exception of the leaf nodes there are a plurality of branches connected to each node. Each branch includes a series connected switch which is used to complete the connection from the node from which it branched to the node connecting said branch to its sub-tree or leaf node. Each of the branching connections at a given node is made by closing a switch in the desired branch and opening a corresponding switch in each of the unselected branches coupled to the node in question. Each of the leaf nodes is connected to two or more "pull-down switches" which connect the non-selected word lines to the inactive power terminal which is usually ground. As will be explained in more detail below, the number of pull-down switches depends on the number of levels of branching in the tree structure.

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)